Aesop Fable for Network Loops

Marc Mosko, Glenn Scott, and Dave Oran

arXiv:1612.04430v1 [cs.NI] 13 Dec 2016

***Abstract*—Detecting loops in data networks usually involves counting down a hop limit or caching data at each hop to detect a cycle. Using a hop limit means that the origin of a packet must know the maximum distance a packet could travel without loops. It also means a loop is not detected until it travels that maximum distance, even if that is many loops. Caching a packet signature at each hop, such as a hash or nonce, could require large amounts of memory at every hop because that cached information must persist for as long as a loop could forward packets. This paper presents a new distributed loop detection mechanism based on a Tortoise and Hare algorithm that can quickly detect loops without caching per-packet data at each hop with a modest amount of additional state in each packet.**

I. Introduction

Detecting loops in data networks usually involves counting down a hop limit or caching data at each hop to detect a cycle. Using a hop limit means that the origin of a packet must know the maximum distance a packet could travel without loops. It also means a loop is not detected until it travels that maximum distance, even if that is many loops. Caching a packet signature at each hop, such as a hash or nonce, could require large amounts of memory at every hop because that cached information must persist for as long as a loop could forward packets. This paper presents a new distributed loop detection mechanism based on a Tortoise and Hare algorithm that can quickly detect loops without caching per-packet data at each hop with a modest amount of additional state in each packet.

One class of algorithms for detecting loops in a series are based on the tortoise and hare construction, attributed to Floyd by Knuth [1]. The core idea is that the tortoise moves one-by-one through a series, denoted as x_i, and the hare moves twice as fast, denoted as x_{2i}. If there is a cycle, eventually $x_{2i} = x_i$, where hop $2i$, modulo the cycle length, is the same item as i. At each step of the algorithm, $tortoise \leftarrow \text{next}(tortoise)$ and $hare \leftarrow \text{next}(\text{next}(hare))$.

Brent's Algorithm [2] finds cycles like Floyd's algorithm, however it requires less memory than Floyd's algorithm to compute the "next" operation. Brent's algorithm compares x_{2^i-1} with all x_j in the open interval $j = [2^i, 2^{i+1})$. Therefore, the same x_{2^i-1} is used at each iteration until there's a single update. $tortoise \leftarrow x_{2^i-1}$, $hare \leftarrow \text{next}(hare)$, until the next power of two when $tortoise \leftarrow hare$. This formation keeps the state update using "next" proceeding in-step with the algorithm iteration, so a packet does not need to carry any memory with it beyond $hare$.

Section II reviews Brent's centralized algorithm. Section III presents our distributed version of the algorithm for use in a packet network.

M. Mosko and G. Scott are with Palo Alto Research Center (PARC)

Algorithm 1 Brent's Algorithm (for base 2)

```
1:  hare ← x_0
2:  power ← 1
3:  hops ← 0
4:  repeat
5:     tortoise ← hare
6:     power ← power * 2
7:     repeat
8:        hops ← hops + 1
9:        hare ← next(hare)
10:    until tortoise = hare or hops ≥ power or hare = ∅
11: until tortoise = hare or hare = ∅
12: if tortoise = hare then
13:    return true
14: end if
15: return false
```

II. Brent's Algorithm

Algorithm 1 shows Brent's algorithm for the case where $\mathcal{Q} = 2$ and $\mu = 0$, which he calls the common case, based on [2]. We have used more descriptive names than x, y, k, r. Mapping Brent's variables to more descriptive names, we use $x \rightarrow hare$, $y \rightarrow tortoise$, $r \rightarrow power$, and $k \rightarrow hops$.

In this formulation, the algorithm tracks four variables: $tortoise$, $hare$, $power$, and $hops$. For a centralized algorithm, tracking this information is not a burden, but in a distributed algorithm we wish to minimize the message sizes to reduce overhead. In our distributed version of the algorithm, in Section III, we reduce the necessary message size to only $tortoise$ and $hops$.

III. Distributed Brent's Algorithm

Our use of Brent's algorithm is shown in Alg 2 and Alg. 3. When a node creates a new packet, it initializes the $tortoise$ variable to its node identifier (nodeid). The node identifier could be the hash a a system's public key or an administratively assigned unique identifier. When a node receives a packet, it first increments the hop count. If the packet's $tortoise$ is equal to the current node's $nodeid$, then there is a loop (the algorithm returns true). Otherwise, if the hop count is equal to the $power$, then the packet has traveled a power of 2 hops, so the algorithm updates the $tortoise$ value.

Based on the hop count, one can detect if the packet has gone a power of 2 distance using the equation 1, a commonly known bit manipulation, where the symbol $\&$ means bit-wise AND.

$$\text{hops is power of 2} \iff (hops \ \& \ (hops - 1)) = 0 \qquad (1)$$

Algorithm 2 Initialize Packet

1: $packet.tortoise \leftarrow nodeid$
2: $packet.hops \leftarrow 0$

Algorithm 3 Receive Packet

1: $packet.hops \leftarrow packet.hops + 1$
2: **if** $packet.tortoise = nodeid$ **then**
3: **return** true
4: **else if** $packet.hops$ is power of 2 **then**
5: $packet.tortoise \leftarrow nodeid$
6: **end if**
7: **return** false

The size of $tortoise$ must be large enough that there is a very low probability of duplicates along a given path. Either the size must be so large that there's a vanishingly small probability that there are no duplicates or we should ensure that a retransmission will not have the same failure mode. Fig. 1 uses the Birthday Paradox to calculate the probability that two or more routers on a path of Path Length have the same NodeId assuming the IDs are picked randomly. For example, for a 32-bit nodeid, there's about a 1% chance of collision for a path of 8192 hops. For paths of practical maximum lengths, such as at most 256 to 512 hops, node ids of 48 to 64 bits should be sufficient.

If we wish to make sure that a retransmission breaks a path with duplicate node IDs, we can need to use a virtual node ID. One could calculate the virtual node ID as

$$id = \text{hash}(trueid, \text{hash}(packet))$$

, where the $trueid$ is a large number, such as the SHA-256 hash of a system's public key, and the hash of the packet includes a short nonce per retransmission and does not include the loop-prevention fields (tortoise, power, hops). This method also better preservers anonymity because the $tortoise$ identity is scrambled, potentially with a cryptographic-grade hash.

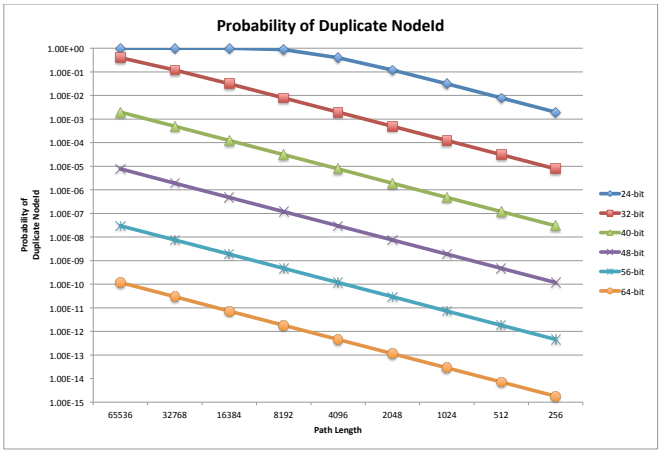

Fig. 1. Probability of duplicates on a path

REFERENCES

[1] D. Knuth, *The Art of Computer Programming*. Addison-Wesley, 1969, vol. 2.
[2] R. Brent, "An improved monte carlo factorization algorithm," *BIT Numerical Mathematics*, vol. 20, no. 2, pp. 176–184, 1980. [Online]. Available: http://dx.doi.org/10.1007/BF01933190